

Authorisation

The present invention relates to authorisation, in particular to authorisation using a digital certificate.

5 It is known to use digital certificates to authorise a node or a user, for example to authorise access to data or services for the node or user. Such certificates normally have a digital signature which is encrypted using a public key cryptography system. The digital signature is normally a function of the characters forming the message content of the certificate, such that a recipient can perform a function on the signature in order to
10 determine with some degree of certainty that a received certificate has not been altered. The use of a digital certificate in authorising data transfer to or from a mobile node is particularly important.

According to one aspect of the invention, there is provided a method of authorising data transfer to or from a mobile node temporarily connected to an attachment
15 point of a network, the attachment point having a forwarding node associated therewith for forwarding messages to or from the mobile node, the method including the steps of: (a) receiving a digital certificate from the forwarding node, which certificate includes a message body and a digital signature for verifying the content of the message body, the message body having geographical information therein, which geographical information is
20 derived from a physical location; (b) performing a comparison between the geographical information of the certificate and a further item of geographical information; and, (c) making an authorisation decision for data transfer to or from the mobile node in dependence on the result of the comparison.

By including geographical information in a certificate from the forwarding node,
25 the likelihood is reduced that an authorisation will be made in error or as a result of fraudulent activity.

Preferably, the geographical information in the certificate will be derived from the physical location of the forwarding node. This will allow the location of the mobile node to be inferred from the geographical information in the certificate of the forwarding node.
30 Location base services and other data can then be provided to the mobile node.

Further aspects of the invention are provided as specified in the appended claims. The present invention is described in further detail below, by way of example only, with reference to the following drawings in which:

Figure 1 shows a network system according to the present invention;

35 Figure 2 is a schematic representation of a digital certificate;

Figure 3 is a schematic representation of message flows between nodes;

Figure 4 shows the transfer of messages involved in the creation of a security association; and,

Figure 5 is a more detailed example of a certificate.

5 In Figure 1, there is shown a network system 10 having a main network 12 and at least one mobile node 14. The main network, which is preferably static, has a plurality of nodes 16 connected by links 18. Each node has an address, the addresses of the main network being arranged in a hierarchical system, such that the address of a node will normally indicate the topological position of that node. In the present example, th
10 addresses of the nodes are addressed according to the Internet Protocol, preferably version V6.

The mobile node 14 is configured to make a temporary connection with any one of a plurality of spaced apart attachment points 20 of the main network 12. Each attachment point will normally have a node, termed a foreign agent (FA) node 22
15 associated therewith (only one is shown for clarity). The foreign agent will normally issue the mobile node with a temporary address, which address is topologically related to that of the issuing foreign agent, (for example, the addresses may share a common prefix portion) such that packets addressed to the temporary or "care-of" address of the mobile node will be routed by the network to the foreign agent, which can then forward the
20 packets to the mobile nodes. The foreign agent will also serve to forward messages from the mobile node to another destination in the main network.

The mobile node has an associated Home Agent (HA) node in the main network 12. The association between the mobile node and its home agent is formed at least in part by a permanent address allocated by the home agent to the mobile node, which the
25 mobile node retains as it moves from one attachment point to another. The permanent or "home" address of the mobile node will be topologically related to the address of the home agent (for example by sharing a common prefix portion with the home agent address) such that packets from a caller node 26 (CN) addressed to the home address of the mobile node can be intercepted by the home agent. To allow the home agent to forward a
30 packet from the caller node 26 towards the current attachment point of the mobile node, the home agent will store a mapping between the current care-of address of the mobile node and its home address, which mapping will be updated when the mobile node attaches to a new attachment point: that is, when the mobile node transmits a binding update to its home agent informing the home agent of its new care-of address.

The mobile node may be a router or a communications device on a vehicle, or otherwise the mobile node may be a portable device, such as a laptop computer, or another type of movable device. Preferably, the mobile node will have temporary connection means 32 for making a temporary connection 34 with an attachment point, for example a radio receiver and/or transmitter for making a radio connection 34, or a releasable electrical or optical connector arrangement.

There are many circumstances in which authentication or other authorisation will be desirable before secure communication between two nodes is established. For example, the home agent for the mobile node may require proof of the identity of the mobile node before accepting a binding update, so as to reduce the risk of traffic intended for the mobile node being inadvertently forwarded by the home agent to a fraudulent node. The need for efficient security processes is particularly important in the case of traffic relating to a mobile node, since the topologically correct address of a mobile node is temporary, that is, changeable as the mobile node moves. However, there are other situations where authorisation or authentication can be important: for example, the home agent may have a policy of only passing information to specified foreign agents, or likewise, a foreign agent may have a policy of only allowing mobile nodes to attach to it whose identity or other characteristics fall within a specified or predetermined category.

To reduce the risk of fraudulent authentication or authorisation, or other data transfer taking place, the main network 12 will normally include a certificate authority agent, here implemented as a certificate authority (CA) node 28. (It will be appreciated that the nodes CA, FA, MN, and HA will be implemented on hardware which will include at least one memory and at least one processor means, the hardware and software running thereon being located at a single node or otherwise being distributed over spaced apart apparatus, for example over a plurality of nodes).

The certificate authority will normally employ a Public Key (PKI) encryption system. In such a system, also known as asymmetric key cryptography, an entity (such as a person or node) has associated therewith a pair of keys: a public key which is publicly accessible, for example by being distributed or being placed in a public directory; and, a private key; only accessible to the entity with which the pair of keys is associated. The pair of keys is mathematically linked, for example according to a known protocol developed by Diffie and Hellman. The mathematical function relating the two keys to one another is such that it is difficult, preferably unfeasible, to derive the private key from the related public key. This may be achieved by a function which requires an impractically large number to be factored in order to obtain the private key. Thus, a first person wishing

to send an encrypted message for transmission to a second person can look up the public key associated with the first person in a public (and trusted) directory, and encrypt the message with the second person's public key. The second person can use their private key to decrypt the message. In this way, public key cryptography can be considered to be
5 based on a one-way function, that is, a function which is significantly easier to perform in the forward direction than in the reverse direction. The public key provides an indication of an instance of the function, and the private key allows the function to be performed in the reverse direction.

In order to generate a certificate, the certificate authority will form a digital
10 signature in association with the information content of the certificate. The digital signature will be the result of a mathematical algorithm, function, or other computation having as input parameters (a) the message content of the certificate, and (b) the private key of the certificate authority. In particular, the digital signature will preferably be the result of the encryption procedure using the private key of the certificate authority. A
15 person wishing to read the certificate can then "de-crypt" the digital signature using the public key of the certificate authority (or equivalently, perform a function to generate signature information related to the encrypted signature). A checking algorithm can then be performed using the digital signature and the message content as input parameters to determine whether the received message corresponds to the digital signature, in particular
20 whether the received message is the same as the transmitted message used to generate the digital signature. This is possible because for a given private key, the digital signature is (almost) unique to the message: that is, the likelihood of two different (non-identical messages) returning the same (even unencrypted) signature is very low. Furthermore, because the digital signature is encrypted, it is difficult for an unauthorised person to
25 change the signature so as to reflect any changes that unauthorised person may have made to the message. In this way, the digital signature is indicative of the message content such that, by performing predetermined respective functions on the received message content and the digital signature, and by comparing the results of those functions, it is possible to determine if the message content as received has been altered.

30 In more detail, to generate a certificate, the certificate authority will: perform a "hash" function on the certificate (message) content, or other function chosen such that there is a low likelihood of two different contents yielding the same result. The result of the hash function, known as the message digest, is then encrypted using the certificate authority's private key according to a PKI protocol. A recipient can then perform a
35 recipient computation, related to that used to create the signature. The recipient

computation function involves the message content, the received signature, and the sender's (here the certificate authority) signature. If the result is correct according to a predetermined mathematical relation, the signature can be deemed genuine, since the message content is unlikely to have been altered.

5 Thus, a recipient can: "de-encrypt" the signature using the public key, in order to obtain the message digest (or related information); perform the same hash function on the received message as was performed on the sent message; and, compare one message digest with the other. If these are the same, the signature is deemed genuine. When an entity (the issuee) is issued a certificate by the certificate authority, the message content
10 of the certificate will normally contain at least some of the following items of information: name of issuing certificate authority; the public key of the certificate authority; an expiration date of the public key; the name or an identifier of the issuee; and, the public key of the issuee.

 In addition, the message content will include location object identifier, or other
15 geographical information, which geographical information is derived from a physical geographical position or an indication thereof. Examples of geographical information include; a latitude and longitude value (with optionally an altitude value); a map reference; a known place name; a street or road name; and, a street junction. Since geographical information is derived from a geographical location, it will be more reliable as an indication
20 of position than other information such as an IP address, from which geographical position can sometimes be inferred.

 A certificate 50 is illustrated in Figure 2, which shows: the message content 52; items of information such as geographical information 54; an identifier 56; and, the digital signature 58.

25 When the certificate authority issues a certificate to a node, the certificate authority will transmit the certificate to the requesting node over the network 12; that is, through one or more routing node 161 and links 18. The requesting node can then store the certificate in a memory, preferably in a local memory 30, such that the certificate can be transmitted to another node when needed, for example when information or services
30 are required from that node.

 Returning to the situation shown in Figure 1, the mobile node 14 and the foreign agent 22 will each be issued with a certificate by the certificate authority 28. The certificate for the mobile node will normally have geographical information indicative of an area associated with the home agent's physical location, but the geographical information
35 in the mobile nodes certificate may be other static geographical information, for example

information relating to the owner's place of residence. In more detail, the geographical information will normally be in the form of a value associated with a location object identifier. Likewise, the home agent and foreign agent are also sent respective certificates by the certificate authority 28. The value for the location object identifier for the home agent and foreign agent correspond to their respective physical locations as expressed in latitude and longitude. In this example, the value of the location object identifier for the mobile node corresponds to that of the home agent.

The steps involved in the attachment of the mobile node to the main network are shown schematically in Figure 3, in which information flow is indicated by arrows, increasing time being in the downward direction on the page. To begin the attachment process, the mobile node sends an initial registration packet to the foreign agent, which packet is "dropped" or read at the foreign agent. The initial packet triggers the start of an Internet Key Exchange (IKE) process for establishing a security association between the mobile node and the foreign agent, in which process protocols are agreed. Once a secure association has been established, the mobile node may send encrypted traffic to the foreign agent.

As part of the registration process between the mobile node and the foreign agent, the mobile node will send its certificate to the foreign agent. The foreign agent can then: de-crypt the digital signature using the public key of the certificate authority, which public key the foreign agent may obtain from the certificate authority itself; perform a function on the content, which function (normally a hash function) has previously agreed (for example during the IKE procedure on the message); compare the result of the function with the decrypted signature; and, if the comparison indicate a match, treat the certificate as genuine. Assuming the certificate is genuine (or to ascertain or further verify that the certificate is genuine), the foreign agent can then extract the location object identifier from the message content of the certificate. The foreign agent may be configured to make a decision as to whether to grant or refuse foreign agent functionality to a mobile node in dependence on the geographical information in the mobile nodes certificate. In particular, the foreign agent may be configured to compare the location object identifier of the mobile node to information indicative of the foreign agent's own physical location information, which may be stored locally, and only grant access if the two items of location information have a specified characteristic in common. For example, access may only be granted if the location information of the mobile node and foreign agent indicate respective positions within the same specified geographical area or within a specified distance of one another. In this way, the foreign agent can be configured to only

grant access to a mobile node which originates from the same geographical district or country as the foreign agent.

After it has been established that the foreign agent can grant access, or provide other foreign agent functionality for the mobile node, the foreign agent will attempt to register with the home agent. To start this process, the foreign agent will transmit an initial registration packet, which packet is "dropped" at the home agent. This dropped packet initiates an IKE procedure as indicated in Figure 4. The home agent will receive a certificate from the foreign agent, and perform similar steps to those outlined above to determine whether the certificate is genuine. That is, the home agent will extract the location object identifier from the certificate of the foreign agent and will perform a comparison between the location object identifier and other stored geographical information. In particular, the home agent may compare the location object identifier against an expected location object identifier stored at a registry 36, which registry may store location object identifiers respectively mapped to the identity of mobile nodes and foreign agent nodes. Thus, the location object identifier in the certificate may serve to provide an additional security test in order to authenticate the foreign agent.

Once the mobile node is registered with the foreign agent, and the foreign agent is registered with the home agent, a secure association is formed on the one hand between the mobile node and the foreign agent, and on the other between the foreign agent and the home agent. Encrypted traffic can then be transmitted from the mobile node to the foreign agent, and then forwarded by the foreign agent to the home agent. However, in some embodiments, only the registration of the foreign agent with the home agent is needed for the mobile node to receive data from the home agent.

As the certificate is used for authentication in this secure association creation process, the location information contained in the certificate, and the associated IP address can be extracted and stored for future use. The home agent will normally have a security policy that grants or denies mobile IP services in dependence upon the location of the foreign agent. When a request for a mobile IP registration arrives at the home agent, the home agent may use the IP address of the request message to obtain the location of the care-of address for the mobile node. However, the certificate from the foreign agent will preferably be used to obtain the physical location of the foreign agent (or a confirmation thereof), as this is more reliable. Once the location of the foreign agent has been obtained, it can be compared against a policy associated with that location. If the mobile node is allowed mobile IP services from the location of the foreign agent (the location of the mobile node being inferred from that of the foreign agent), then a

registration-successful message will be sent back to the foreign agent, else a registration-unsuccessful message will be sent back.

It can be seen from the above that the location information in a certificate can be used by a node when deciding whether to provide information. In particular, the location
5 information extracted from a certificate can be compared with stored location information, such that the decision as to whether surfaces are to be provided can be made at least in part in dependence upon the comparison between the extracted location information and the stored location information.

By using the certificate from a foreign agent (forwarding node) to make an
10 authorisation decision, advantage can be taken of the increase security associated with a fixed node over that associated with a mobile node.

Further details on the implementation of one embodiment of the invention are provided below: the operating used is FreeBSD [FreeBSD]. The Internet Key Exchange (IKE) implementation comes from KAME [kame], the secure socket layer implementation
15 comes from the openssl organisation [openssl] and the mobile IP implementation from Portland State University [psu]. The openssl code is used by the KAME IKE implementation. It is also assumed that security policy exists that state that secure communication must exist between the MN and the FA and also between the FA and the HA. One stage is to introduce the location attribute into the certificate. This is done by
20 introducing a new object identifier of type 2.5.5.4 [oid] and associating a value with this corresponding to the location expressed as x, y pair. It is also possible to include an altitude attribute as x, y, z where z represents the altitude although this was not done in this implementation. An example of such a certificate is shown in Figure 5 with the location object identifier and associated value shown underlined.

Another stage is to configure the IKE daemon, racoon [kane], to use certificates
25 rather than pre-shared secrets. As shown schematically in Figure 3, the sending of the registration packet from the MN to FA initiates the generation of a security association between them. Lets focus on the creation of security associations between the FA and the HA since the creation of security associations between the FA and the MN is as
30 described in standards [RFC2002]. Figure 4 shows the sequence of messages that occur in phase 1 of the creation of secure associations using IKE. Note that in Figure 4, the certificate payload has to be present since it may not be possible for the FA and the HA to get the certificate from other sources, say secure DNS. Where the diagram ends, phase 2 of the IKE processing can take place to create the IPsec secure association proper. It is
35 intended to send the valid certificate to a local listener that will store the location and the

IP address in a local file. The message from the IKE daemon is parsed. With regard to the processing of the certificate payload: the function saveFaLocation (currentLocation, ip_address) saves the location seen in the certificate and the associated ip address as a tuple in an ascii file. This file can then be read by other applications that require location dependent information. The home agent may have a policy for allowing mobility, which could be refined by defining bounded polygons for location, as in [RFC2009].

References

- [RFC1712] <http://www.ietf.org/rfc/rfc1712.txt>
- 10 [geobytes] <http://www.geobytes.com>
- [RFC2002] <http://www.ief.org/rfc/rfc2002.txt>
- [newbury] <http://www.newburynetworks.com>
- [RFC2401] <http://www.ietf.org/rfc/rfc2401.txt>
- [FreeBSD] <http://www.freebsd.org>
- 15 [kame] <http://www.kame.net>
- [openssl] <http://www.openssl.org>
- [psu] <http://www.cs.pdx.edu/research/SMN/index.html>
- [oid] <http://www.alvestrand.no/objectid/>